

How to sell an app: pay-per-play or buy-it-now?

Shuchi Chawla^{*} Nikhil R. Devanur[†] Anna R. Karlin[‡] Balasubramanian Sivan[§]

Abstract

We consider pricing in settings where a consumer discovers his value for a good only as he uses it, and the value evolves with each use. We explore simple and natural pricing strategies for a seller in this setting, under the assumption that the seller knows the distribution from which the consumer's initial value is drawn, as well as the stochastic process that governs the evolution of the value with each use.

We consider the differences between up-front or “buy-it-now” pricing (BIN), and “pay-per-play” (PPP) pricing, where the consumer is charged per use. Our results show that PPP pricing can be a very effective mechanism for price discrimination, and thereby can increase seller revenue. But it can also be advantageous to the buyers, as a way of mitigating risk. Indeed, this mitigation of risk can yield a larger pool of buyers. We also show that the practice of offering free trials is largely beneficial.

We consider two different stochastic processes for how the buyer's value evolves: In the first, the key random variable is how long the consumer remains interested in the product. In the second process, the consumer's value evolves according to a random walk or Brownian motion with reflection at 1, and absorption at 0.

^{*}University of Wisconsin-Madison. shuchi@cs.wisc.edu.

[†]Microsoft Research. nikdev@microsoft.com.

[‡]University of Washington. karlin@cs.washington.edu.

[§]Microsoft Research. bsivan@microsoft.com.

1 Introduction

A standard assumption in mechanism design about a consumer of goods is that he has valuations for various bundles of goods, which he knows ahead of time. The consumer is then modeled as a utility maximizer, where the utility is typically quasi-linear – equal to the consumer’s valuation minus payments made by him. While this is often a reasonable assumption, in this paper we propose an alternate model that is more appropriate for many scenarios. In particular, we consider scenarios where the consumer discovers his valuation of a good as he *uses* it. For instance, consider a consumer buying a song on iTunes. The consumer’s valuation of the song depends on how much he enjoys it as he listens to it and how many times he wants to listen to it. The consumer does not know these quantities ahead of time, and only discovers them as he repeatedly listens to the song. Indeed the consumer’s enjoyment of the song, rather than being constant over time, may evolve as he listens to the song more and more. Another example is an app or a video game that a consumer plays repeatedly. As in the case of a song, the consumer’s value may evolve as he uses the product repeatedly; At some point of time, the consumer may tire of the product altogether and stop using it.

We model such scenarios by having a value per usage that evolves with usage. We use V_t to denote the value to the consumer for the $t + 1^{\text{th}}$ usage¹. The consumer knows V_t only after she has used the good for t times. Note that t denotes usage and not real time: if the consumer does not use the item his valuation for the next usage does not change. We model the evolution of V_t as a random process. The consumer has an initial value for the very first usage, V_0 , which is sampled from a distribution. Subsequently, V_t evolves according to a given random process that depends on V_0 .

The prevalent mechanism for selling songs, apps, software, and other digital goods is to offer a one-time contract to the consumer for unlimited usage of the product. We call such a mechanism a “Buy-It-Now” (BIN) scheme. Our value evolution model can be “reduced” back to the standard valuation model for a one-shot game by considering the expected sum of all the V_t s conditioned on V_0 ; This is the total expected value that the consumer obtains through unlimited use of the product². However, such a reduction overlooks the fact that there are other interesting pricing mechanisms in the value evolution model. For instance, a simple pricing mechanism is to charge the consumer *per usage*. We call such a mechanism a “Pay-Per-Play” (PPP) scheme. While any one-shot mechanism offers a uniform price for unlimited usage to all consumers, a PPP scheme allows different consumer types to pay different amounts based on how long they stay interested in the product. It therefore price-discriminates in a natural and effective way. A PPP scheme also gives the consumer a finer grained control over his own utility, and is therefore preferable to consumers that are risk averse.

Another example of a mechanism in the value evolution model is the common practice of offering a “free-trial” to the consumers. Following this free trial, consumers that stay interested may purchase the product in one shot or per-usage. Our model offers a formal explanation of the benefits of such a free trial period.

We would like to highlight the imminently practical nature of the model we consider. Almost none of the goods such as music, games, or apps are currently sold on a PPP scheme, yet such a scheme would be extremely easy to implement given current technology. Insights obtained using our model could therefore be influential in the adoption of a new pricing scheme for such goods. Our initial findings suggest that different PPP schemes could do better (and sometimes *much* better) than the standard BIN pricing scheme. Similarly,

¹In some cases, it may be appropriate to model usage as a continuous quantity, such as the duration of time a game is played, but for the following discussion it is easier to think of it as discrete, and this assumption does not affect our analysis.

²This assumes that the consumer is risk neutral, which is perhaps *not* a reasonable assumption in the relevant scenarios. In case of risk averse consumers, one can likewise consider the risk adjusted expectation of total value.

our results suggest that offering a free trial is often a very good idea. Yet, these findings only scratch the surface and they open up a whole new spectrum of open problems and research directions (which we discuss towards the end of the paper).

We consider two different random models of value evolution, both of which are meant to be idealized versions of value evolution in some relevant markets. The first model is a binary value model where the consumer's value stays constant while he continues to remain interested in the product, and then drops to zero. In other words, there is a joint distribution over V_0 and another random variable T , and $V_t = V_0$ for $t \leq T$ and $V_t = 0$ otherwise. The binary value model captures scenarios where the consumer's experience from using the product stays the same over time, but after a few uses, the consumer either does not require the product any more or switches to a competitor. Software usage is a good example of this. The other model we study is a Martingale random walk model where V_t is Markovian and the increments are independent with mean 0.³ The random walk model captures the notion that the value evolves continuously with usage, each usage either increasing or decreasing the value by a small amount. We emphasize that the details of the models are not the focus of the paper; they are just meant to be tools for formal analysis by which we may confirm our intuition. Quantitatively, our results depend on the nature of the random process that governs value evolution, but qualitatively we draw insights that apply to both models.

We now summarize our results for the different value evolution models and risk attitudes. Our first (and strongest) result is for the binary value model, and the rest are for the random walk model.

- **PPP versus BIN in the binary value model.** We show that there exists a PPP scheme whose revenue is higher than that of any BIN scheme even if the consumer is infinitely risk averse⁴ in responding to the PPP pricing, and risk neutral in responding to the BIN scheme. Furthermore, the PPP scheme generates more consumer utility than the BIN scheme. Thus, the PPP scheme is better than BIN in all respects: seller revenue, consumer utility and consumer risk. (Theorem 3.1)
- **Price discrimination.** We formalize the intuition that PPP schemes are good price discriminators, by showing that for a risk neutral consumer, there exists a PPP scheme in which *each* consumer type, given by V_0 , pays at least half of his *cumulative value* in expectation, where the cumulative value is his total value if he keeps using the good⁵. This is *almost perfect price discrimination*. (Perfect price discrimination would be each consumer paying his cumulative value.) Thus, a single PPP scheme obtains revenue at least $\frac{1}{2}$ of social welfare, for any initial distribution on V_0 . (Theorem 4.6)
- **Free trial.** We show that free trial can also lead to good price discrimination: appropriately choosing a trial period (number of free usages) and then offering a buy-it-now price to a risk neutral consumer also has each consumer pay a constant fraction of his cumulative value in expectation. (Theorem 4.8)
Furthermore, a combination of **free trial and PPP** enables price discrimination even over infinitely risk averse consumers, and gives the strongest result in the random walk model: infinitely risk averse consumers pay a constant fraction of their cumulative value under this scheme. (Theorem 4.7)
- **Risk aversion.** Intuitively, a PPP scheme is better for a risk averse consumer, since his payment can depend on the V_t s and he therefore has much better control over his realized utility. At the extreme end of the spectrum, where the consumer is infinitely risk averse, we show that a PPP pricing scheme

³The model is quite general and includes as a special case the Brownian motion, a standard model of the evolution of stock prices. See Appendix B for details.

⁴See Section 2 for a definition of infinite risk aversion.

⁵It also turns out that the cumulative value itself is proportional to V_0 .

can get an asymptotically⁶ unbounded factor higher revenue than BIN. We also give an interpolation of this result for a single parameter family of risk averse utilities that range from the infinitely risk averse to risk neutral. (Theorems 4.9 and A.1)

- The fact that the risk neutral PPP revenue is half of the social welfare also implies that it is half of the BIN revenue, but which is higher? The answer actually depends on the initial distribution. Given any distribution, we show how to compute the two revenues so they can be compared. While we don't have a concise characterization of distributions for which one of the schemes is better, our computations for a few distributions confirm the following intuition: distributions that are not heavily concentrated towards the upper end of $[0, 1]$ favor the PPP pricing scheme. Many of the well known distributions have this property, and for the uniform distribution we obtain that a PPP scheme gets a higher revenue than the optimal BIN scheme, while having each consumer type pay (in expectation) less than the optimal buy-it-now price. In other words, the PPP scheme grows the total pie and hence gets a bigger share of the pie. (Remarks 4.3 and 4.4)

Together these results give us a better understanding of the strengths and weaknesses of the different pricing schemes. They are witnesses to the kind of insights we can derive using our model of value evolution. This is yet simply scratching the surface and we expect in future many more interesting results in this model.

Related Work. Initiated by the work of Baron and Besanko [3], there have been a number of papers [6, 10, 4, 11, 16, 5] that consider revenue maximizing contracts where the private information of the agents evolves over time. The main difference from our model is that in this line of work, the agents' private information evolves over time *irrespective* of whether the agent consumes the good or not. In contrast, in our model the buyer's value evolves only if he consumes the good. The similarity to our model is that the value is often modeled as evolving according to a Markov process. We note that both models are applicable in different scenarios: pure evolution with time is probably more applicable for a gym membership whereas ours is more applicable for products such as music, software, or video games. To our knowledge, the evolution of consumers' values with usage has not been studied previously in mechanism design literature.

A technically unrelated direction, falling under the broad umbrella of online mechanisms, studies revenue/welfare maximization when buyers arrive sequentially, with either adversarial value distributions, or values drawn from a known/unknown distribution [13, 2, 9, 12, 8, 1]. A special kind of online mechanism is a dynamic pricing mechanism where the seller posts a price for each buyer (as opposed to more general schemes like auctions). There is a huge body of work on dynamic pricing and revenue management problems in the Operations Research literature. See [7] and references there in. The book by Vohra and Krishnamurthi [17] compares various pricing strategies like posted-prices, auctions, and haggling in the presence/absence of competition.

2 Model & Preliminaries

Basic Setup. We consider a single seller offering goods for repeated private consumption (e.g. digital goods), and a single buyer. We normalize the seller's cost of production to 0. The buyer's initial private value for the good, $V_0 \in [0, 1]$, is drawn from a publicly known distribution F . The buyer's value for the good evolves with repeated consumption of the good. He knows the process by which V_t evolves but only

⁶The asymptotics are with the variance going to zero, which is due to the choice that we bound the value per usage to be within $[0, 1]$. Since the quantities are scale invariant, we "scale up" by decreasing the variance.

learns his actual value V_t for the $(t + 1)^{st}$ consumption immediately after consuming the good t times. The seller, on the other hand, only knows the underlying distributions: the initial value distribution F and the process by which the V_t evolves, but not its precise value. We assume that the buyer's value remains bounded in the range $[0, 1]$ and is non-zero for only a finite number of steps. The buyer faces uncertainty as to how his value will evolve in the future. This uncertainty exposes him to some risk, and he takes this into account while making his purchase decisions. We elaborate further on buyer behavior below. The seller's goal is to maximize his expected revenue, which is the sum total of all the prices paid by the buyer.

Buyer's value evolution. We consider two models for how the buyer's value evolves.

- **The binary value model:** In the binary value model, the buyer's value remains at his initial value V_0 for some number of steps $T(V_0)$, and then falls to 0 in the $(T(V_0) + 1)$ th step. The number of steps $T(V_0)$ is unknown to both the buyer and the seller, but it is drawn from a publicly known distribution. This model captures settings where the buyer makes a binary decision to continue or discontinue the use of the product, but as long as he continues, he has a fixed value for the product or service. For example, the buyer may derive some fixed amount of pleasure from playing a video game or going to the gym but switches to a competitor after some rounds of play. We assume that $E[T(V_0)|V_0]$ is non-decreasing in V_0 and is finite.
- **The random walk model:** In the random walk model after each consumption, the buyer's value increases or decreases by some fixed $\delta > 0$, with probability $\frac{1}{2}$ each. When the value reaches 0, it gets absorbed there and doesn't change anymore. In other words, the buyer loses interest in the product and does not want to continue purchasing it at any price. When the value hits 1, it gets reflected, i.e., it decreases by δ . Both the buyer and the seller are aware of this process of value evolution.

The random walk model is a special case of a more general *Markov evolution* model, in which the additive change to the buyer's value after the t th consumption is a random variable, Δ_t , that depends only on t . We assume that Δ_t has mean 0 and standard deviation $\delta > 0$ (assumed to be small).

These models capture settings where each consumption can slightly enhance or diminish the buyer's overall experience of the product. Magazine subscriptions, subscriptions to episodes of a TV serial, or to a radio channel, are some examples.

Selling mechanisms. Perhaps the most prevalent selling mechanism for digital goods such as songs, apps, and video games, is what we call the *buy-it-now* (BIN) mechanism. Here the seller charges a one-time fee for unlimited usage of the product. We compare BIN mechanisms to schemes that charge the buyer for each successive consumption of the product. We collectively call such schemes *pay-per-play* (PPP) mechanisms. A pay-per-play scheme is specified by an infinite vector of prices \mathbf{p} , with p_t denoting the price charged for the t th consumption of the product. In this paper we focus on simple PPP schemes such as a constant price PPP scheme. Pay-per-play schemes found in practice, such as magazine subscriptions and gym memberships, typically charge a fixed amount per usage.

We also analyze the effect of offering a *free trial*: the seller offers the product for free for some fixed number of steps, and thereafter either offers a buy-it-now price, or a constant price PPP scheme.

Buyer cumulative value, social welfare, and utility. The cumulative value of a buyer is the total value he can potentially accumulate before he loses interest in the product (his value drops to 0). Formally, consider a buyer with initial value V_0 whose value evolves according to one of the models above. Let $T(V_0)$ be the

buyer's stopping time – the time at which his value drops to 0 or below. Then the *cumulative value* of this buyer is the random variable $\sum_{0 \leq t < T(V_0)} V_t$. We denote by $\mathcal{C}(v)$ the expectation of this random variable conditioned on $V_0 = v$.

The *social welfare* generated by a mechanism is the total value that the buyer obtains from using the product. In a BIN setting, the buyer's *social welfare* is his cumulative value if he purchases the good, and 0 otherwise. In a PPP setting, if we denote by X_t the indicator random variable which is 1 if the buyer purchases at the t^{th} opportunity, then the buyer's social welfare with a starting value of V_0 is $\sum_{t \geq 1} V_{t-1} X_t$. Finally, the *utility* of a buyer with initial value V_0 in a BIN setting with price p is $\sum_{0 \leq t < T(V_0)} V_t - p$ if the buyer purchases, and 0 otherwise. In a PPP setting, the *future utility* of a buyer that has purchased t times, and whose current value is V_t , facing price sequence \mathbf{p} is $\sum_{s \geq t} (V_s - p_{s+1}) X_s$.

Buyer's purchase behavior and risk aversion. In a buy-it-now mechanism, the buyer makes a one-time decision to purchase or not purchase the product, depending on his (random) future utility. In a pay-per-play mechanism, the buyer makes a purchasing decision at every step of the process, once again taking into account the effect of his decision on his (random) future utility. Because the value of the buyer as well as the price of the product evolve with consumption, once the buyer decides to discontinue buying and consuming the product, the value and price freeze, and the buyer never buys the product again.

We assume that the buyer is a utility maximizer, but because of the uncertainty in his future value for the item, is also sensitive to losses. In the main body of the paper, we consider the two extreme ends of the spectrum with respect to risk: risk neutral buyers on the one hand, and completely (or infinitely) risk averse buyers on the other. A *risk neutral* buyer chooses to buy or not depending on which of these maximizes his expected utility. An *infinitely risk-averse* buyer not only ensures that his future utility is maximized in expectation, but also ensures that his future utility will be non-negative with probability 1. We use $\mathcal{R}_0^{\text{BIN}}$ (resp. $\mathcal{R}_0^{\text{PPP}}$) to denote the revenue from a BIN scheme (resp. PPP scheme) with a risk-neutral buyer. $\mathcal{R}_\infty^{\text{BIN}}$ and $\mathcal{R}_\infty^{\text{PPP}}$ denote the corresponding revenues from an infinitely risk-averse buyer.

In Appendix A, we consider a single-parameter family of risk-averse preferences that interpolates between risk neutral and infinitely risk averse preferences.

Monopoly price and single-round optimal revenue. We will need the optimal revenue achievable in a single round when a buyer's value is drawn from the publicly known distribution F . A special case of Myerson's result [15] shows that the optimal revenue in this case is $\max_p p(1 - F(p))$. The price that optimizes this expression is the monopoly price, denoted by μ , is given by $\phi^{-1}(0)$, where $\phi(x) = x - \frac{1-F(x)}{f(x)}$, and the optimal revenue is denoted by $\mathcal{R}^{\mathcal{M}}$.

3 Binary value model

We begin with an analysis of the binary value model of value evolution. Recall that in this model, the buyer has a fixed constant value for usage of the product. However, after some random number of steps, $T(V_0)$, which is unknown to both the buyer and the seller, the buyer loses interest in the product and his value drops to 0. We call $T(V_0)$ the buyer's stopping time. Since the buyer's value doesn't change as long as he maintains interest in the product, a simple and natural selling mechanism is a constant price PPP scheme. Regardless of the buyer's risk attitude, a buyer will purchase the product in the constant price scheme if and only if V_0 exceeds the price, and in that case, for as long as his value is non-zero. Buyers with different stopping times pay different amounts to the mechanism, which gives an effective way to price discriminate.

We convert this intuition into a proof, showing that a constant price PPP scheme with an appropriate price obtains more revenue (even with a risk averse buyer) than a BIN scheme (even with a risk neutral buyer). Moreover, the PPP scheme generates more social welfare because it serves more buyers. Some of this social welfare can be shared with the buyers, resulting in a higher total utility for the buyers in addition to a higher revenue for the seller. So in this setting, PPP is better in all respects than BIN.

Theorem 3.1. *In the binary value model when $\mathbf{E}[T(V_0)|V_0]$ is a non-decreasing function of V_0 there is a constant price PPP scheme that with an infinitely risk averse buyer gets at least as much revenue, as much social welfare, and as much buyer utility as the revenue optimal BIN scheme with a risk-neutral buyer.*

Proof. In the binary value model with a *risk-neutral* buyer, recall that the buyer's expected cumulative value is $\mathcal{C}(V_0) = V_0 \mathbf{E}[T(V_0)|V_0]$. The optimal BIN price is the monopoly price for the distribution of this random variable where V_0 is drawn from the distribution F . We assume that $\mathbf{E}[T(V_0)|V_0]$ is a non-decreasing function of V_0 . So there is a threshold initial value v^{BIN} such that the risk neutral buyer purchases in the BIN scheme if and only if his initial value is at least v^{BIN} . The revenue of the BIN scheme is then $\mathcal{R}_0^{\text{BIN}} = \mathcal{C}(v^{\text{BIN}})(1 - F(v^{\text{BIN}})) = v^{\text{BIN}} \mathbf{E}[T(v^{\text{BIN}})](1 - F(v^{\text{BIN}}))$.

Now consider a constant price PPP scheme. We begin by setting the per-play price in the scheme to be equal to the threshold value v^{BIN} . At this price, any buyer with $V_0 \geq v^{\text{BIN}}$ purchases the product until his value becomes 0. This is the same set of buyers that purchase in the BIN scheme described above. Moreover, a buyer with initial value V_0 pays a total of $v^{\text{BIN}} \mathbf{E}[T(V_0)]$, which is at least as large as $v^{\text{BIN}} \mathbf{E}[T(v^{\text{BIN}})]$, the amount the same buyer pays in BIN. Thus, $\mathcal{R}_\infty^{\text{PPP}} \geq \mathcal{R}_0^{\text{BIN}}$. Further, since this PPP scheme serves the same set of buyers as the BIN scheme, the two generate the same total social welfare. Now, consider gradually decreasing the per-play price in the PPP scheme to below v^{BIN} . Then, the social welfare generated by PPP increases, while its revenue may or may not decrease. We can continue decreasing the price as long as the revenue of PPP stays above that of BIN, and at some point, both the revenue and the buyer utility (which is social welfare – revenue) of PPP exceed the corresponding quantities for BIN. ■

4 The random walk model

We now switch to the random walk model of value evolution. Recall that in this model, the buyer's value evolves as a random walk with step size $\delta > 0$. The value stays bounded within $[0, 1]$; When it reaches 1, it gets reflected back to stay within the range $[0, 1]$. When it reaches 0, it gets absorbed, or remains 0. In Appendix B we consider a more general Markov evolution model and show that many of our results extend to that setting.

4.1 Notation and basic facts

Let $h_{u,v}$ denote the expected time for a random walk starting at u to hit v , and, $h_{u,\{v,v'\}}$ denote the expected time for a random walk starting at u to hit one of v or v' . We use $\mathbf{P}_v[E]$ to denote the probability of a (random-walk-related) event E when the random walk starts at v . Similarly $\mathbf{E}_v[X]$ denotes the expected value of a (random-walk-related) random variable X when the random walk starts at v .

The following lemma is standard. See, e.g., [14]. Note that this lemma is about a simple random walks with step size δ (i.e., there is no reflection at 1 or absorption at 0).

Lemma 4.1. Let X_t be a simple random walk with increments of $\pm\delta$, starting from $X_0 = v$ with $v \in (0, 1)$. Let $w < v < u$, Then,

$$\mathbf{P}_v [X_t \text{ reaches } u \text{ before } w] = \frac{v - w}{u - w} \quad (1)$$

In particular,

$$\mathbf{P}_v [X_t \text{ reaches } 1 \text{ before } 0] = v$$

and,

$$h_{v, \{w, u\}} = \frac{(v - w)(u - v)}{\delta^2} \quad (2)$$

The following lemma summarizes several important facts about our random walk process. In particular, it determines the hitting times for this walk taking reflection at 1 and absorption at 0 into account. Some of the analysis later in the paper is based on the fact that buyers with high initial values are likely to quickly reach the maximum value. This is formalized in the third part of the following lemma, which derives the expected time to reach 1 conditioned on reaching 1 before 0.

Lemma 4.2. Suppose that V_t evolves according to the random walk model starting at $V_0 = v$. Then

$$\mathbf{P}_v [V_t \text{ reaches } 1 \text{ before } 0] = v. \quad (3)$$

For any $v > u$

$$h_{vu} = \mathbf{E}_v [\text{time to hit } u] = \frac{(v - u)(2 - v - u)}{\delta^2} \implies T(v) = \mathbf{E}_v [\text{time to hit } 0] = \frac{v(2 - v)}{\delta^2}. \quad (4)$$

Let τ be the first time V_t hits 0 or 1. Then,

$$\mathbf{E}_v [\tau | V_\tau = 1] = \frac{1 - v^2}{3\delta^2}. \quad (5)$$

Proof. Equation (3) is a standard fact about this kind of random walk. Equation (4) follows from Lemma 4.1, Equation (2), and the reflection principle: the random walk with reflection at 1 and absorption at 0 is equivalent to a random walk between $[0, 2]$ with absorption at both 0 and 2. Thus, reaching u in the old walk is equivalent to reaching one of u and its reflection $2 - u$ in the new walk.

Equation (5) is Exercise 17.1 in [14] (proved by considering a martingale related to V_t , namely $M_t = Y_t^3 - 3tY_t$, where $Y_t = \frac{V_t}{\delta}$). (See also Appendix B.) ■

Expected cumulative value. Next we present expressions for the buyer's cumulative value. Let $\mathcal{C}(v, w)$ denote the expected cumulative value of a buyer with starting value v , until the value reaches w for the first time.

Lemma 4.3. Let $v > w$. Then,

$$\mathcal{C}(v, w) = \frac{v - w}{\delta^2} (v(1 - v) + (1 - w)(1 - \delta) + \delta^2)$$

Proof. Using Lemmas 4.1 and 4.2, we have that the expected cumulative value of a buyer with starting value v , until the value reaches w for the first time is:

$$\begin{aligned}\mathcal{C}(v, w) &= v \cdot \mathbf{E}_v [\text{Time to hit } w \text{ or } 1] + \mathbf{P}_v [\text{Hitting 1 before } w] \cdot \mathcal{C}(1, w) \\ &= v \cdot \frac{(v - w)(1 - v)}{\delta^2} + \frac{v - w}{1 - w} \mathcal{C}(1, w)\end{aligned}$$

Similarly:

$$\begin{aligned}\mathcal{C}(1, w) &= 1 + \mathcal{C}(1 - \delta, w) \\ &= 1 + \frac{(1 - \delta)(1 - \delta - w)\delta}{\delta^2} + \frac{1 - \delta - w}{1 - w} \mathcal{C}(1, w) \\ \Rightarrow \mathcal{C}(1, w) &= (1 - w) \left[\frac{1}{\delta} + \frac{(1 - \delta)(1 - \delta - w)}{\delta^2} \right]\end{aligned}$$

Plugging the expression for $\mathcal{C}(1, w)$ into that for $\mathcal{C}(v, w)$, we obtain the result. ■

Corollary 4.4. *The expected cumulative value of a buyer with $V_0 = v$ is:*

$$\mathcal{C}(v) = \mathcal{C}(v, 0) = \frac{v^2(1 - v)}{\delta^2} + \frac{v}{\delta^2} (1 - \delta + \delta^2) \in \frac{v}{\delta^2} [1, 5/4].$$

Optimal risk-neutral BIN revenue. Suppose that $V_0 = v$. A risk-neutral BIN buyer will compute his expected cumulative value with this initial value, namely, $\mathcal{C}(v)$, and accept a price of p if and only if $\mathcal{C}(v) \geq p$. Thus the risk neutral BIN revenue is just the Myerson optimal revenue for the distribution of $\mathcal{C}(V_0)$, and the optimal price is the monopoly price for the distribution of $\mathcal{C}(V_0)$. Thus we have:

Lemma 4.5. *The optimal risk-neutral BIN revenue satisfies $\frac{\mathcal{R}^{\mathcal{M}}}{\delta^2} \leq \mathcal{R}_0^{\text{BIN}} \leq \frac{5\mathcal{R}^{\mathcal{M}}}{4\delta^2}$ where $\mathcal{R}^{\mathcal{M}}$ is the Myerson optimal revenue in one round for the distribution F .*

4.2 PPP enables near-perfect price discrimination over risk-neutral buyers

In this section we will show a surprising fact about PPP: if the buyer is risk-neutral, PPP can enable near-perfect price discrimination and obtain a constant fraction of the buyer's expected cumulative value. The PPP scheme achieving this bound is extremely simple: we charge a constant price of $1/2$ to the buyer for each usage of the product.

Theorem 4.6. *Consider a constant price PPP scheme that charges a per-usage price of $\frac{1}{2}$. The risk-neutral revenue of this scheme is at least half of the expected cumulative value of the buyer:*

$$\forall v \in [0, 1], \quad \mathcal{R}_0^{\text{PPP}}|_{V_0=v} \geq \frac{1}{2} \mathcal{C}(v).$$

Proof. How does a risk-neutral PPP buyer react to a constant price PPP at a price of p ? At any point of time, this buyer computes his expected future utility conditioned on buying at this step, and buys if this expectation is non-negative. Because the PPP price stays constant, the buyer's strategy is time-independent and has a threshold behavior: the buyer continues buying until his value drops below a threshold w . We will now compute this w as a function of the PPP price p . Let $U(v, w, p)$ denote the buyer's expected future

utility in a constant price PPP scheme with per-usage price p when the buyer's current value is v and he continues buying until his value drops below w . The buyer's threshold is then the smallest w such that for all $v \geq w$, $U(v, w, p)$ is non-negative.

Recall that $\mathcal{C}(v, w)$ denotes the expected cumulative value of a buyer with starting value v , until the value reaches w for the first time. We can use Lemma 4.3 to compute $U(v, w, p)$.

$$\begin{aligned} U(v, w, p) &= \mathcal{C}(v, w) - p \cdot \mathbf{E}_v [\text{Time to hit } w] = \mathcal{C}(v, w) - p \cdot \frac{(v-w)(2-w-v)}{\delta^2} \\ &= \frac{v(v-w)(1-v)}{\delta^2} + \frac{(v-w)(1-w)(1-\delta)}{\delta^2} + (v-w) - p \cdot \frac{(v-w)(2-w-v)}{\delta^2} \\ &= \frac{v-w}{\delta^2} (1-\delta-w(1-\delta-p) + v(1-v) - p(2-v) + \delta^2) \end{aligned} \quad (6)$$

The RHS of equation (6) is a decreasing function of w . Further, as long as $v \geq w$ and $p \leq \frac{1}{2}$, the RHS is non-negative. Therefore, at a constant price of $\frac{1}{2}$, a risk-neutral buyer keeps buying until his value hits 0. Surprisingly, even at a value of δ , the buyer is willing to pay a price of $\frac{1}{2}$ because he expects to eventually make up for the lost utility. This means that $\mathcal{R}_0^{\text{PPP}}|_{V_0=v} = \frac{1}{2}T(v)$. On the other hand, $\mathcal{C}(v) \leq T(v)$, since the per usage values are always at most 1. The lemma follows. ■

Remark 4.1. It appears counter-intuitive at first glance that with a price of $\frac{1}{2}$ the buyer keeps buying until his value hits 0. The intuition is that although the buyer may be losing money at some particular moment in time, either his value will eventually get high, and he will have large gains for a relatively long time, or his value will go to 0 relatively quickly. The resulting expected value is positive.

Remark 4.2. Theorem 4.6 shows that posting a price of $\frac{1}{2}$ ensures that the buyer continues buying until his value hits 0. What if the price is larger than $\frac{1}{2}$? Then the buyer's threshold to stop buying turns out to be $w = 2p - 1$. To see this, consider a buyer at value $v = w + \delta$; If his value goes one step lower he will stop buying. Such a buyer's utility is $U(w + \delta, w, p) = \frac{1}{\delta} [(1-w-\delta)(1+w-2p) + \delta(1-p)]$. It follows that $U(w + \delta, w, p) \geq 0$ if $w \geq 2p - 1$.

Remark 4.3. Note that Theorem 4.6 gives an exact expression for the payment of a consumer with initial value v in the $\frac{1}{2}$ price PPP scheme, namely $\frac{1}{2}T(v) = \frac{v(2-v)}{2\delta^2}$. Similarly from Corollary 4.4, we have an exact expression for $\mathcal{C}(v)$. Thus, for any given distribution, we can compute the optimal BIN revenue and the constant price PPP revenue with a price of $\frac{1}{2}$, and compare the two to see which is higher.

Remark 4.4. Theorem 4.6 guarantees that PPP obtains at least half of the revenue of BIN for risk neutral buyers, however, this comparison is often quite loose. In fact, for many distributions over the initial value, $\mathcal{R}_0^{\text{PPP}} \geq \mathcal{R}_0^{\text{BIN}}$. For instance, when V_0 is drawn from the $U[0, 1]$ distribution, the expected revenue of PPP is $\frac{1}{2\delta^2} \int_0^1 v(2-v)dv = \frac{1}{3\delta^2}$. On the other hand, $\mathcal{R}_0^{\text{BIN}}$ can be computed as follows. We have $\mathcal{C}(v, 0) \cong \frac{v^2(1-v)}{\delta^2} + \frac{v}{\delta^2}$. The BIN optimal revenue is obtained by computing $\max_t \mathcal{C}(t, 0)(1 - F(t))$. Taking $F(t) = t$, the maximum value for this expression can be computed to be $\frac{5}{16\delta^2} < \frac{1}{3\delta^2}$. Moreover, the optimal BIN price is $> \frac{1}{2\delta^2}$ (corresponding to a threshold initial value of $\frac{1}{2}$), while the expected payment of a consumer with value v in PPP is always $\leq \frac{1}{2\delta^2}$.

Remark 4.5. On the other hand, there are distributions for which the revenue of any constant price PPP scheme is strictly smaller than $\mathcal{R}_0^{\text{BIN}}$. For instance when $F(x) = x^2$, a constant price PPP with a price of $1/2$ obtains revenue $\frac{5}{12\delta^2}$; Using the discussion in remark 4.2, it is straightforward to show that prices larger or smaller than $\frac{1}{2}$ perform no better. On the other hand, we can compute the revenue of the optimal BIN scheme as $\mathcal{R}_0^{\text{BIN}} = \max_t \mathcal{C}(t, 0)(1 - F(t)) = \frac{0.479}{\delta^2}$.

4.3 Free trial enables near-perfect price discrimination over all risk profiles

We now show that near-perfect price discrimination can be achieved even over risk averse buyers, if the pricing scheme allows the buyer to try the product for free for a carefully chosen number of trials. We first discuss a PPP scheme with a free trial period that achieves this. Next we show that even BIN with a free trial period enables near-perfect price discrimination over risk neutral agents.

In the following lemmas, let $\mathcal{R}^{\text{free+PPP}}(T, c)$ be the expected revenue of the seller obtained by offering the buyer a free trial period of length T , and then offering him a pay-per-play price of c . Likewise, let $\mathcal{R}^{\text{free+BIN}}(T, c)$ be the expected revenue of the seller obtained by offering the buyer a *free trial period* of length T , and then offering him a buy-it-now price of c .

Theorem 4.7. *Let $T = \frac{2}{3\delta^2}$ and $c = 0.089$. Then the expected seller revenue from a PPP pricing scheme with a free trial period of length T and PPP price c for a risk averse buyer is a constant fraction of the buyer's expected cumulative value:*

$$\forall v \in [0, 1], \quad \mathcal{R}_{\infty}^{\text{free+PPP}}(T, c)|_{V_0=v} = \Theta(\mathcal{C}(v)).$$

Proof. By Equation (5) in Lemma 4.2, and Markov's inequality with $T = \frac{2}{3\delta^2}$,

$$\mathbf{P}_v \left[\tau > T \mid V_{\tau} = 1 \right] \leq \frac{1}{2}. \quad (7)$$

Let T' be the length of time after the free trial period ends for which V_t is at least c : $T' = \min\{t - T : t > T, \text{ and } V_{t+1} < c\}$. Then, $\mathcal{R}_{\infty}^{\text{free+PPP}} = \mathbf{E}[cT']$. Let τ denote the time at which the buyer's value reaches 0 or 1. Then for a buyer with starting value $V_0 = v$ we obtain

$$\begin{aligned} \mathcal{R}_{\infty}^{\text{free+PPP}}(T, c) &\geq \mathbf{E}_v [cT' \mid V_{\tau} = 1 \text{ and } \tau \leq T] \mathbf{P}_v [\tau \leq T \mid V_{\tau} = 1] \mathbf{P}_v [V_{\tau} = 1] \\ &\geq c \mathbf{E} [h_{1c} - T] \frac{1}{2} v \quad (\text{by Inequality (7) and Equation (1) in Lemma 4.1}) \\ &\geq c \cdot \frac{(1-c)^2 - 2/3}{\delta^2} \cdot \frac{v}{2} \quad (\text{by Equation (4) in Lemma 4.2}) \end{aligned}$$

This latter quantity for sufficiently small c is $\Omega\left(\frac{v}{\delta^2}\right)$. It is maximized at $c \approx 0.089$. Corollary 4.4 then implies the lemma. ■

Theorem 4.8. *Let $T = \frac{3}{8\delta^2}$ and $c = \frac{1}{4\delta^2}$. Then the expected seller revenue from offering a free trial period of length T and then a BIN price of c to a risk neutral buyer is a constant fraction of the buyer's expected cumulative value:*

$$\forall v \in [0, 1], \quad \mathcal{R}_0^{\text{free+BIN}}(T, c)|_{V_0=v} = \Theta(\mathcal{C}(v)).$$

Proof. Suppose that the value of the risk neutral buyer at the end of the trial period, V_T is at least $c\delta^2$. Then, noting that $\mathcal{C}(x) \geq \frac{x}{\delta^2}$, the buyer accepts a BIN price of c . We now lower bound the probability $Q(v)$ that $V_T \geq c\delta^2$ given $V_0 = v$. We can write $Q(v) \geq Q_1(v)Q_2$, where

$$Q_1(v) = \mathbf{P} [V_{\tau} = 1 \text{ and } \tau \leq T \mid V_0 = v] \quad \text{and} \quad Q_2 = \mathbf{P} [V_T \geq c\delta^2 \mid V_{\tau} = 1 \text{ and } \tau \leq T].$$

By Equation (1) in Lemma 4.1, Equation (5) in Lemma 4.2, and Markov's inequality, with $T = \frac{3}{8\delta^2}$, we have $Q_1(v) \geq \frac{v}{9}$.

Next, recall that if $X = \sum_{i=1}^n X_i$ with each X_i being ± 1 with probability half each, independently, then $\mathbf{P}[|X| \geq a] \leq e^{-\frac{a^2}{2n}}$. We can bound Q_2 from below by setting $n = T - \tau$, and X_i to be $1/\delta$ times the change in the buyer's value at step $\tau + i$, and requiring that the net change in value from step τ to step T is no more than $(1 - c\delta^2)$. Thus,

$$Q_2 \geq 1 - \mathbf{P}\left[|X| \geq \frac{1 - c\delta^2}{\delta}\right] \geq 1 - 2e^{-\frac{(1 - c\delta^2)^2}{2(T - \tau)\delta^2}} \geq 1 - 2e^{-3/4} = 0.055$$

Clearly, $\mathcal{R}_0^{\text{free}+\text{BIN}} \geq Q(v) \cdot c \geq \frac{0.055 \cdot v}{9} \cdot \frac{1}{4\delta^2} = \Theta(\mathcal{C}(v))$. ■

Remark 4.6. An interesting fact about Theorems 4.7 and 4.8 is that the proposed pricing does not depend on the distribution of the buyer's initial value. This demonstrates that in the random walk model the evolution of the value is a more important phenomenon than the initial value at which the walk starts.

4.4 Comparing BIN and PPP

Theorems 4.6 and 4.7 compare the PPP revenue to the buyer's cumulative value under various risk profiles. As a direct consequence, PPP gets a large fraction of the revenue of any BIN scheme as well. In this section, we make the comparison between PPP and BIN more precise and show that in some cases PPP beats BIN by far. The following theorem deals with the case of risk averse buyers. (In Appendix A we extend this result to the case of α -risk averse buyers for some $\alpha \in [0, \infty)$ achieving a tradeoff between PPP and BIN revenue as a function of the risk parameter.)

Theorem 4.9. *With an infinitely risk averse buyer, there exists a constant price PPP scheme with revenue $\mathcal{R}_\infty^{\text{PPP}} \geq \frac{1}{4\delta} \mathcal{R}_\infty^{\text{BIN}}$.*

Proof. Revenue of BIN. In BIN, the infinitely risk-averse buyer will first compute his cumulative value under the worst possible evolution of values, namely, when his value goes down by δ in every single round, and thus, will hit 0 in v/δ rounds. The buyer's worst possible cumulative value is then $\frac{v^2}{2\delta}$. From this it follows that $\mathcal{R}_\infty^{\text{BIN}} = \max_t \frac{t^2}{2\delta} (1 - F(t))$.

Revenue of PPP. Let $t^* = \arg\max_t \frac{t^2}{2\delta} (1 - F(t))$. Consider a constant price PPP scheme with a price of $\frac{t^*}{2}$ per round. We get the revenue from an infinitely risk averse buyer to be:

$$\begin{aligned} \mathcal{R}_\infty^{\text{PPP}} &= \frac{t^*}{2} \text{ times the expectation over } V_0 \text{ of the time to hit } t^*/2 \\ &= \frac{t^*}{2} \frac{1}{\delta^2} \int_{\frac{t^*}{2}}^1 \left(v - \frac{t^*}{2}\right) \left(2 - v - \frac{t^*}{2}\right) f(v) dv \\ &\geq \frac{(t^*)^2}{4\delta^2} \int_{t^*}^1 \left(2 - v - \frac{t^*}{2}\right) f(v) dv \\ &\geq \frac{(t^*)^2}{8\delta^2} \int_{t^*}^1 f(v) dv = \frac{(t^*)^2}{8\delta^2} (1 - F(t^*)) = \frac{1}{4\delta} \mathcal{R}_\infty^{\text{BIN}} \end{aligned}$$

Here the first inequality follows by restricting the limit of the integral to values greater than t^* and the second follows by noting $v, t^* \leq 1$. ■

Remark 4.7. The factor $\Theta(\frac{1}{\delta})$ is tight. Consider a point mass distribution at v . Here, we have $\mathcal{R}_\infty^{\text{BIN}} = \frac{v^2}{2\delta}$, while $\mathcal{R}_\infty^{\text{PPP}} = \max_p p \cdot \frac{(v-p)(2-p-v)}{\delta^2} \leq \frac{v^2}{\delta^2}$.

Next we show that in some settings, PPP with a constant price obtains revenue from a risk averse buyer that compares well even to the BIN revenue from a risk neutral buyer.

Theorem 4.10. *There exists a PPP scheme, that, with an infinitely risk-averse buyer gets a revenue at least a $\frac{\mu}{10}$ fraction of the revenue optimal BIN scheme with a risk-neutral buyer, where μ is the monopoly price of the distribution F of the initial value V_0 .*

Proof. Recall that by Lemma 4.5, the optimal risk-neutral BIN revenue satisfies $\mathcal{R}_0^{\text{BIN}} \leq \frac{5\mathcal{R}^\mathcal{M}}{4\delta^2}$.

We now consider the infinitely risk-averse PPP revenue for a constant per-play price p . This revenue is given by $\mathcal{R}_\infty^{\text{PPP}} = p \cdot \int_p^1 h_{vp} f(v) dv$, where h_{vp} , as before, is the time taken by the random walk starting at v to go strictly below p . By Equation (4) in Lemma 4.2, we have $h_{vp} = \frac{(v-p)(2-p-v)}{\delta^2}$. Now consider a price of $p = \frac{\mu}{2}$. Since $p < 1/2$, we have $2 - p - v \geq \frac{1}{2}$. Thus, as in the proof of the previous lemma, we have,

$$\mathcal{R}_\infty^{\text{PPP}} \geq \frac{p}{2\delta^2} \int_p^1 (v-p)f(v)dv \geq \frac{\mu}{4\delta^2} \frac{\mu}{2} \int_\mu^1 f(v)dv = \frac{\mu}{8\delta^2} \mathcal{R}^\mathcal{M} \geq \frac{\mu}{10} \mathcal{R}_0^{\text{BIN}}.$$

■

Remark 4.8. When the monopoly price μ of the initial value distribution F is a constant, PPP with a risk-averse buyer gets a constant fraction of BIN with a risk-neutral buyer. Otherwise, PPP gets a very small fraction of the BIN revenue. This is not just an artifact of our analysis: suppose that the buyer's initial value is v with probability 1. Then BIN gets a revenue of $\Omega(v/\delta^2)$. On the other hand, PPP with a constant price of p can obtain at most $2p(v-p)/\delta^2$, which is a factor of $\Omega(v)$ smaller than the $\mathcal{R}_0^{\text{BIN}}$.

5 Conclusions

In this paper, we have studied pricing in settings where a consumer learns his value for a good only as he uses it. We have seen that in such a setting, the use of “pay-per-play” pricing can be advantageous to both the seller, in terms of revenue, and to the buyer, as a mechanism for reducing risk. Likewise, offering a free trial period can improve the seller's revenue.

The results in this paper are just first steps. Some natural future questions include the following: (1) What is the optimal seller mechanism and how does it depend on how risk-averse the buyer is? As an intermediate model between PPP and BIN pricing, it may help to price usage at regular intervals, as a sort of “subscription” pricing. (2) Within the framework of PPP pricing, is it a good idea to set a constant price per usage, or does it help to increase or decrease price over time? (3) What about other models for value evolution? For example, it would be interesting to incorporate drift, either positive or negative, into the buyer's value evolution. Likewise, can we say something interesting when per-step changes in value are not always small? (4) Can we generalize these results to incorporate other models of risk? (5) What happens when multiple schemes co-exist, e.g. when the consumer is offered both a BIN price and a PPP scheme?

It would be very interesting to also have empirical estimation of consumer behavior to complement the theoretical results. Experiments that shed light on what is a good model of value evolution would be extremely useful. Similarly, risk aversion models that explain consumer behavior in these markets would add great value.

References

- [1] Moshe Babaioff, Shaddin Dughmi, Robert Kleinberg, and Aleksandrs Slivkins. Dynamic pricing with limited supply. In *ACM Conference on Electronic Commerce, EC '12, Valencia, Spain, June 4-8, 2012*, pages 74–91, 2012.
- [2] Ziv Bar-Yossef, Kirsten Hildrum, and Felix Wu. Incentive-compatible online auctions for digital goods. In *Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '02*, pages 964–970, 2002.
- [3] David P. Baron and David Besanko. Regulation and information in a continuing relationship. *Information Economics and Policy*, 1(3):267 – 302, 1984.
- [4] Marco Battaglini. Long-term contracting with markovian consumers. *American Economic Review*, 95(3):637–658, 2005.
- [5] Dirk Bergemann and Philipp Strack. Dynamic revenue maximization: A continuous time approach. 2014.
- [6] David Besanko. Multi-period contracts between principal and agent with adverse selection. *Economics Letters*, 17(12):33 – 37, 1985.
- [7] Omar Besbes and Assaf Zeevi. Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research*, 57(6):1407–1420, November 2009.
- [8] Avrim Blum and Jason D. Hartline. Near-optimal online auctions. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 1156–1163, 2005.
- [9] Avrim Blum, Vijay Kumar, Atri Rudra, and Felix Wu. Online learning in online auctions. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, January 12-14, 2003, Baltimore, Maryland, USA.*, pages 202–204, 2003.
- [10] Pascal Courty and Li Hao. Sequential screening. *Review of Economic Studies*, 67(4):697–717, 2000.
- [11] Peter Eso and Balzs Szentes. Optimal information disclosure in auctions and the handicap auction. *Review of Economic Studies*, 74(3):705–731, 2007.
- [12] Robert D. Kleinberg and Frank Thomson Leighton. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *44th Symposium on Foundations of Computer Science (FOCS)*, pages 594–605, 2003.
- [13] Ron Lavi and Noam Nisan. Competitive analysis of incentive compatible on-line auctions. In *Proceedings of the 2Nd ACM Conference on Electronic Commerce, EC '00*, pages 233–241, 2000.
- [14] David Asher Levin, Yuval Peres, and Elizabeth Lee Wilmer. *Markov chains and mixing times*. American Mathematical Society, 2009.
- [15] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.
- [16] Alessandro Pavan, Ilya Segal, and Juuso Toikka. Dynamic mechanism design: A myersonian approach. *Econometrica*, 82(2):601–653, 2014.

[17] Rakesh V. Vohra and Lakshman Krishnamurthi. *Principles of Pricing - An Analytical Approach*. Cambridge University Press, 2012. ISBN 978-1-107-01065-9.

A α -risk averse BIN vs α -risk averse PPP

We have shown that in the two extreme cases, namely, with risk neutral buyers (Section 4.2) and with infinitely risk averse buyers (Section 4.4) PPP does well compared to BIN. Is it possible to compare the revenue of BIN and the revenue of PPP as a function of how risk-averse the buyer is? We address this question in this section.

α -risk averse buyer. We parameterize risk by the maximum loss the buyer is willing to suffer. Recall that a completely risk averse buyer not only ensures that his future utility is maximized in expectation, but also ensures that his future utility will be non-negative with probability 1. On the other hand, an α -risk averse buyer, for some $\alpha > 0$, is willing to tolerate a loss of up to $1/\alpha$ while maximizing his expected future utility.

Formally, we have:

- **α -risk averse BIN buyer.** Given a BIN price of p , for every $\alpha \geq 0$, an α -risk averse buyer buys if and only if his future utility is non-negative in expectation and no smaller than $-1/\alpha$ with probability 1. Formally, $\mathcal{C}(V_0) - p \geq 0$, and, $\mathbf{P}[\sum_{0 \leq t < T(V_0)} V_t - p < -\frac{1}{\alpha}] = 0$.
- **α -risk averse PPP buyer.** The behavior of the buyer in a PPP mechanism is more complex. As we mentioned earlier, the buyer buys the product for some number of steps and then stops buying it. After the t th step, if the buyer's value V_t is larger than the current price p_t in the PPP scheme, then it is always in the buyer's best interest to continue buying at that step. If V_t is less than p_t , however, then the buyer may still decide to continue buying, banking on the possibility of obtaining a large positive utility later on. Fixing the vector of PPP prices \mathbf{p} , the buyer's strategy specifies as a function of his current value whether or not he should continue to purchase the product. We can define the optimal such strategy via backwards induction. In essence, the optimal PPP buyer strategy maximizes his expected utility, subject to the constraint that his expected utility never falls below $-\frac{1}{\alpha}$. We omit the formal specification of the optimal strategy here because for the price vectors we construct, the optimal strategy is immediate.

Remark A.1. An α -risk averse buyer with $\alpha = 0$ corresponds to a risk neutral buyer, and at $\alpha = \infty$ corresponds to an infinitely risk averse buyer.

We now construct a PPP scheme that obtains a revenue comparable to that of the optimal BIN scheme under any risk profile, that is, any $\alpha > 0$. Our PPP scheme charges a fixed price of p per usage for the first few rounds, and then charges a price of 0 for the remaining rounds. We call such a scheme a *Rent-To-Own* scheme: after the buyer has paid some minimum amount of money to the seller, he gets unlimited use of the product for free. Recall that $\mathcal{R}_\alpha^{\text{PPP}}$ and $\mathcal{R}_\alpha^{\text{BIN}}$ denote the revenues of the optimal PPP and BIN schemes respectively, with an α -risk averse buyer.

Theorem A.1. *Given an α -risk averse buyer with $\alpha \in (0, \infty)$, let v^* be the smallest buyer value that accepts the optimal BIN price. Then,*

$$1. \mathcal{R}_\alpha^{\text{BIN}} \leq \left[\frac{(v^*)^2}{2\delta} + \min\left(\frac{1}{\alpha}, \mathcal{C}(v^*)\right) \right] (1 - F(v^*))$$

2. There exists a PPP scheme with revenue $\mathcal{R}_\alpha^{\text{PPP}} \geq \Theta(1) \cdot \left[\frac{(v^*)^2}{\delta^2} + \min(\frac{1}{\alpha}, \mathcal{C}(v^*)) \right] (1 - F(v^*))$

In particular, for all α , $\mathcal{R}_\alpha^{\text{PPP}} \geq \frac{1}{32} \mathcal{R}_\alpha^{\text{BIN}}$, and when $\frac{1}{\alpha} = O(\frac{(v^*)^2}{2\delta})$, we have $\mathcal{R}_\alpha^{\text{PPP}} = \omega(\mathcal{R}_\alpha^{\text{BIN}})$.

Proof. BIN revenue. If a buyer with value v^* accepts the optimal BIN price, it follows that the optimal BIN price $p^* \leq \mathcal{C}(v^*)$. Also, by the definition of an α -risk averse buyer, the optimal BIN price p^* is at most $\frac{1}{\alpha}$ plus the cumulative value obtained by v^* when the value evolves by decreasing in every round. The latter is the area of a triangle with height v^* and base $\frac{v^*}{\delta}$. Thus, the optimal BIN price $p^* \leq \frac{1}{\alpha} + \frac{(v^*)^2}{2\delta}$. Since we have shown two upper bounds on p^* , it follows that $p^* \leq \min\left(\frac{1}{\alpha} + \frac{(v^*)^2}{2\delta}, \mathcal{C}(v^*)\right) \leq \frac{(v^*)^2}{2\delta} + \min(\frac{1}{\alpha}, \mathcal{C}(v^*))$. Since only buyers with value no smaller than v^* buy, $\mathcal{R}_\alpha^{\text{BIN}} \leq \left[\frac{(v^*)^2}{2\delta} + \min(\frac{1}{\alpha}, \mathcal{C}(v^*)) \right] (1 - F(v^*))$.

PPP revenue. Suppose that among the two terms in the upper bound of p^* , namely $\frac{(v^*)^2}{2\delta}$ and $\min(\frac{1}{\alpha}, \mathcal{C}(v^*))$, we have $\frac{(v^*)^2}{2\delta} \geq \sqrt{\delta} \cdot \min(\frac{1}{\alpha}, \mathcal{C}(v^*))$. In this case the proof of Theorem 4.9 shows that there exists a constant price PPP with revenue $\frac{(v^*)^2}{8\delta^2} (1 - F(v^*)) \geq \Omega(\frac{1}{\sqrt{\delta}}) \mathcal{R}_\alpha^{\text{BIN}}$, proving the theorem, and showing that $\mathcal{R}_\alpha^{\text{PPP}} = \omega(\mathcal{R}_\alpha^{\text{BIN}})$ in this case.

Suppose on the other hand, $\frac{(v^*)^2}{2\delta} < \sqrt{\delta} \min(\frac{1}{\alpha}, \mathcal{C}(v^*))$. Consider a dual price PPP scheme that charges a price of $\min(\frac{1}{2}, \frac{1}{c\alpha\mathcal{C}(v^*)})$ for $c\mathcal{C}(v^*)$ rounds, and 0 thereafter, where $c = 24$. Note that with this PPP scheme, the total possible loss for the buyer is upper bounded by $\frac{1}{c\alpha\mathcal{C}(v^*)} \cdot \mathcal{C}(v^*) \leq \frac{1}{\alpha}$. From the proof of Theorem 4.6 we know that given a price no larger than $\frac{1}{2}$, every buyer type except 0 gets a non-negative future utility by continually buying till the end, where as the buyer is sure to get a 0 utility by rejecting in this round. Since the current round price of $\min(\frac{1}{2}, \frac{1}{c\alpha\mathcal{C}(v^*)}) \leq \frac{1}{2}$, it follows that every buyer type except 0 will buy till the end, resulting in revenue of

$$\mathcal{R}_\alpha^{\text{PPP}} \geq \min\left(\frac{1}{2}, \frac{1}{c\alpha\mathcal{C}(v^*)}\right) \cdot (1 - F(v^*)) \cdot \mathbf{E}_{v^*} [\min(h_{v^*,0}, c\mathcal{C}(v^*))].$$

The term $\mathbf{E}_{v^*} [\min(h_{v^*,0}, c\mathcal{C}(v^*))]$ gives a lower bound on the number of rounds that the seller can collect a price of $\min(\frac{1}{2}, \frac{1}{c\alpha\mathcal{C}(v^*)})$. For the buyer with value v^* , the number of rounds where the seller can collect this price is exactly $\mathbf{E}_{v^*} [\min(h_{v^*,0}, c\mathcal{C}(v^*))]$, and the number of rounds is larger for larger v^* . Given that $\mathbf{E}[h_{v^*,0}] = \frac{v^*(2-v^*)}{\delta^2} \leq \frac{2v^*}{\delta^2}$, and $\mathcal{C}(v^*) \geq \frac{v^*}{\delta^2}$, we have $\mathcal{C}(v^*) \geq \frac{\mathbf{E}[h_{v^*,0}]}{2}$. From Lemma A.2 we know that $\mathbf{P}[h_{v^*,0} \geq k \cdot \mathbf{E}[h_{v^*,0}]] \leq 2^{-\lfloor \frac{k}{2} \rfloor}$. Setting $c = 24$, we get

$$\begin{aligned} \mathbf{E}_{v^*} [\min(h_{v^*,0}, 24\mathcal{C}(v^*))] &\geq \mathbf{E}_{v^*} [\min(h_{v^*,0}, 12 \mathbf{E}[h_{v^*,0}])] \\ &\geq \mathbf{E}[h_{v^*,0}] - \sum_{\ell=6}^{\infty} (2^{-\ell} - 2^{-(\ell+1)}) \cdot (2\ell + 2) \cdot \mathbf{E}[h_{v^*,0}] \\ &= \mathbf{E}[h_{v^*,0}] \cdot \left(1 - 0.5 \cdot \sum_{\ell=6}^{\infty} (2\ell + 2) \cdot 2^{-\ell}\right) \\ &= \frac{3}{4} \mathbf{E}[h_{v^*,0}] \\ &\geq \frac{3}{4} \mathcal{C}(v^*) \end{aligned}$$

If $\min\left(\frac{1}{2}, \frac{1}{24\alpha\mathcal{C}(v^*)}\right) = \frac{1}{2}$, then

$$\mathcal{R}_\alpha^{\text{PPP}} \geq \frac{1}{2} \cdot \frac{3}{4} \cdot \mathcal{C}(v^*)(1 - F(v^*)) = \frac{3}{8} \mathcal{C}(v^*)(1 - F(v^*)) \geq \frac{3}{8} \mathcal{R}_\alpha^{\text{BIN}},$$

where the last inequality follows from $\mathcal{R}_\alpha^{\text{BIN}} \leq \mathcal{C}(v^*)(1 - F(v^*))$.

If $\min\left(\frac{1}{2}, \frac{1}{24\alpha\mathcal{C}(v^*)}\right) = \frac{1}{24\alpha\mathcal{C}(v^*)}$, then

$$\mathcal{R}_\alpha^{\text{PPP}} \geq \frac{1}{24\alpha\mathcal{C}(v^*)} \cdot \frac{3}{4} \cdot \mathcal{C}(v^*)(1 - F(v^*)) \geq \frac{1}{32\alpha}(1 - F(v^*)) \geq \frac{1}{32(1 + \sqrt{\delta})} \mathcal{R}_\alpha^{\text{BIN}},$$

where the last inequality follows from $\mathcal{R}_\alpha^{\text{BIN}} \leq \left(\frac{1}{\alpha} + \frac{(v^*)^2}{2\delta}\right)(1 - F(v^*)) \leq (1 + \sqrt{\delta})\frac{1}{\alpha}(1 - F(v^*))$, given our assumption that $\frac{(v^*)^2}{2\delta} < \sqrt{\delta} \min(\frac{1}{\alpha}, \mathcal{C}(v^*))$. As $\delta \rightarrow 0$, this is a 32 approximation, thus proving the theorem. ■

Lemma A.2. *The time $h_{v,0}$ to hit 0 starting from v satisfies $\mathbf{P}[h_{v,0} \geq k \cdot \mathbf{E}[h_{v,0}]] \leq 2^{-\lfloor \frac{k}{2} \rfloor}$ for $k \geq 2$.*

Proof. From Markov's inequality we have $\mathbf{P}[h_{v,0} \geq 2 \cdot \mathbf{E}[h_{v,0}]] \leq \frac{1}{2}$. By reflection principle, a random walk with reflection at 1 hitting 0 is the same as a random walk without reflection hitting one of 0 or 2. Thus, the random walk with reflection at 1 not having hit 0 after $2 \mathbf{E}[h_{v,0}]$ steps is equivalent to the random walk without reflection not having hit either of 0 or 2 after $2 \mathbf{E}[h_{v,0}]$ steps. The position of the random walk is now a random variable X with expectation v . Our goal is to repeatedly apply Markov's inequality. To do this, we need an upper bound on $\mathbf{E}[h_{X,0}]$. We have $\mathbf{E}[h_{X,0}] = \frac{2\mathbf{E}[X] - \mathbf{E}[X^2]}{\delta^2} \leq \frac{2\mathbf{E}[X] - \mathbf{E}[X]^2}{\delta^2}$. Substituting $\mathbf{E}[X] = v$, we get $\mathbf{E}[h_{X,0}] \leq \frac{2v - v^2}{\delta^2} = \mathbf{E}[h_{v,0}]$. Thus, the probability that, after $2 \mathbf{E}[h_{v,0}]$ more steps the random walk did not hit 0 or 2 incurs an additional factor of at most $\frac{1}{2}$ by Markov's inequality. Applying Markov's inequality $\lfloor \frac{k}{2} \rfloor$ times, we get the lemma. ■

B A more general random walk model

In Section 4.1, we assumed that the value of the buyer evolved according to a simple random walk with a step size of δ , with reflection at 1 and absorption at 0. The same results go through assuming the value evolves according to Brownian motion with mean 0 and standard deviation δ . In particular, the necessary lemmas used in the proofs, e.g., all the results in Section 4.1, hold essentially unchanged. (See the book by Levin, Peres and Wilmer [14], Theorem 2.49, Page 57 and Exercise 2.17 (b), page 62.)

In this section, we consider a more general discrete time random walk model, and outline some of the (standard) arguments used to prove the claims in Section 4.1. Specifically, we assume that the value V_t evolves according to a discrete time Markov process, which is a martingale (except at the boundaries). We can also forego the assumption that the walk is symmetric (implying the third moment of the steps is 0). We will assume though that: (a) $V_0 \in (0, 1)$. (b) Reflection occurs whenever V_t crosses 1 (from below). Specifically, if $V_t > 1$, then $V_{t+1} := V_{t-1}$. Note that in a departure from the basic random walk model, we will allow the value to temporarily exceed 1 to simplify our analysis. (c) There is absorption (i.e., the buyer loses interest forever), whenever V_t drops to 0 or below. (d) There is a constant $\epsilon > 0$ which is an upper bound on $\Delta_t := |V_{t+1} - V_t|$ for every t . In what follows we assume that ϵ is sufficiently small.

Finally, we use the notation:

$$\delta^2 := \mathbf{E}[(\Delta_t)^2 | V_t] \quad \text{and} \quad c_3 := \mathbf{E}[(\Delta_t)^3 | V_t].$$

Definition 1. Define the stopping time τ to be the first time at which $V_t \geq 1$ or $V_t \leq 0$.

Our theorems in Section 4 follow from the following facts that are generalizations of Lemmas 4.1 and 4.2, and Corollary 4.4.

1. $\mathbf{P}_v[V_t \text{ hits } 1 \text{ before } 0] \in v \pm O(\epsilon)$. (Equation (3) in Lemma 4.2)

2. Let τ be the stopping time as defined above. Then (special case of Equation (2) in Lemma 4.1)

$$\mathbf{E}_v[\tau] \in \frac{v(1-v) \pm O(\epsilon)}{\delta^2}.$$

3. Let $w = V_{\tau+1}$. Then for any $x < w$, (special case of Equation (4) in Lemma 4.2)

$$h_{w,x}|_{V_\tau \geq 1} = \mathbf{E}_w[\text{time to hit } x \text{ or less} | V_\tau \geq 1] = \frac{(1-x)^2 \pm O(\epsilon)}{\delta^2}.$$

4. The expected cumulative value starting from v is (Corollary 4.4)

$$\mathcal{C}(v) \in \Theta\left(\frac{v \pm \epsilon}{\delta^2}\right).$$

5. The expected time to reach a value of 1 or higher conditioned on reaching that value before 0 is (special case of Equation (5) in Lemma 4.2)

$$\mathbf{E}_v[\tau | V_\tau \geq 1] = \frac{[1 - v^2 - c_3(1-v)](1 \pm O(\epsilon))}{3\delta^2}.$$

The proofs of the above facts follow from standard martingale arguments that we outline here.

For (1): τ is a stopping time and V_t is a martingale in $[0, \tau - 1]$, so by the optional stopping theorem

$$\mathbf{E}_v[V_\tau] = \mathbf{E}_v[V_0] = v = \mathbf{P}_v[V_\tau \geq 1] \mathbf{E}_v[V_\tau | V_\tau \geq 1] + (1 - \mathbf{P}_v[V_\tau \geq 1]) \mathbf{E}_v[V_\tau | V_\tau < 1].$$

which yields this fact since $\mathbf{E}_v[V_\tau | V_\tau \geq 1] \in [1, 1 + \epsilon]$, and, $\mathbf{E}_v[V_\tau | V_\tau < 1] = \mathbf{E}_v[V_\tau | V_\tau \leq 0] \in (-\epsilon, 0]$.

For (2): Use the fact that $X_t = \frac{V_t^2}{\delta^2} - t$ is a martingale, and optional stopping at τ , which yields

$$\frac{v^2}{\delta^2} = \mathbf{E}_v[X_0] = \mathbf{E}_v[X_\tau] = \frac{\mathbf{E}_v[V_\tau^2]}{\delta^2} - \mathbf{E}_v[\tau].$$

An application of (1) completes the argument.

For (3) and (4): follow the proofs given in section 4 (with the appropriate modification for handling the boundary error).

For (5) we slightly generalize Exercise 17.1 from Levin et al. [14].

Let $Y_t = \frac{V_t}{\delta}$. Then, for $t \in [0, \tau - 1]$, Y_t is a martingale. In addition, if we define $c'_3 = (c_3/\delta^3)$, then $M_t = Y_t^3 - 3tY_t - tc'_3$ is also a martingale. To verify this fact, let $\tilde{\Delta}_t = Y_{t+1} - Y_t$.

$$\begin{aligned} M_{t+1} - M_t &= \left[(Y_t + \tilde{\Delta}_t)^3 - 3(t+1)(Y_t + \tilde{\Delta}_t) - (t+1)c'_3 \right] - (Y_t^3 - 3tY_t - tc'_3) \\ &= 3Y_t^2\tilde{\Delta}_t + 3Y_t\tilde{\Delta}_t^2 + \tilde{\Delta}_t^3 - 3t\tilde{\Delta}_t - 3Y_t - 3\tilde{\Delta}_t - c'_3. \end{aligned}$$

Hence

$$\mathbf{E} [M_{t+1} - M_t | V_t] = 3Y_t^2 \mathbf{E} [\tilde{\Delta}_t | V_t] + 3\frac{V_t}{\delta} \mathbf{E} [\tilde{\Delta}_t^2 | V_t] + \mathbf{E} [\tilde{\Delta}_t^3 | V_t] - 3t \mathbf{E} [\tilde{\Delta}_t | V_t] - 3\frac{V_t}{\delta} - c'_3$$

and since $\tilde{\Delta}_t = \frac{\Delta_t}{\delta}$, this equals

$$= 3\frac{V_t}{\delta} + \frac{c_3}{\delta^3} - 3\frac{V_t}{\delta} - c'_3 = 0.$$

Finally, we apply the Optional Stopping Theorem and get

$$\mathbf{E}_v [M_0] = \frac{v^3}{\delta^3} = \mathbf{E}_v [M_\tau] = \mathbf{E}_v \left[\frac{V_\tau^3}{\delta^3} - 3\tau \frac{V_\tau}{\delta} - c'_3 \tau \right]$$

or equivalently

$$v^3 = \mathbf{E}_v [V_\tau^3 - 3\delta^2 \tau V_\tau - c_3 \tau] . \quad (*)$$

Rewrite

$$\begin{aligned} \mathbf{E}_v [V_\tau^3] &= \mathbf{E}_v [V_\tau^3 | V_\tau \geq 1] \mathbf{P}_v [V_\tau \geq 1] + \mathbf{E}_v [V_\tau^3 | V_\tau \leq 0] \mathbf{P}_v [V_\tau \leq 0] \\ &= \mathbf{P}_v [V_\tau \geq 1] (1 + O(\epsilon)) \end{aligned}$$

and

$$\begin{aligned} \mathbf{E}_v [\tau V_\tau] &= \mathbf{E}_v [\tau V_\tau | V_\tau \geq 1] \mathbf{P}_v [V_\tau \geq 1] + \mathbf{E}_v [\tau V_\tau | V_\tau \leq 0] \mathbf{P}_v [V_\tau \leq 0] . \\ &= \mathbf{E}_v [\tau | V_\tau \geq 1] \mathbf{P}_v [V_\tau \geq 1] (1 + O(\epsilon)) \end{aligned}$$

Plugging back into (*), dividing both sides by $\mathbf{P}_v [V_\tau \geq 1] = v(1 \pm O(\epsilon))$ and using fact (2), we obtain

$$v^2 = [1 - 3\delta^2 \mathbf{E}_v [\tau | V_\tau \geq 1] - c_3(1 - v)] (1 \pm O(\epsilon))$$

Rearranging gives the result.

With the above facts in hand, the other results in the paper for the random walk model follow in this more general setting mutatis mutandis.